

Crux: Locality-Preserving Distributed Systems

Michael F. Nowlan, Jose Faleiro, and Bryan Ford
Yale University

Abstract

Distributed systems achieve scalability by balancing load across many machines, but wide-area distribution can introduce worst-case response latencies proportional to the network’s delay diameter. Crux is a general framework to build *locality-preserving* distributed systems, by transforming some existing scalable distributed algorithm \mathcal{A} into a new algorithm $\bar{\mathcal{A}}$ that guarantees for any two clients u and v interacting via service requests to $\bar{\mathcal{A}}$, these interactions exhibit worst-case response latencies proportional to the network delay between u and v . Locality-preserving PlanetLab deployments of a **memcached** distributed cache, a **bamboo** distributed hash table, and a **redis** publish/subscribe service indicate that Crux is effective and applicable to a variety of existing distributed algorithms. Crux achieves several orders of magnitude latency improvement for localized interactions at the cost of increasing per-node overheads, as each physical node must participate in multiple instances of algorithm \mathcal{A} .

1 Introduction

Today’s “big data” algorithms achieve scalability by balancing storage, computation, and communication costs across large distributed clusters, often scattered across data centers around the world. A key challenge in such algorithms is preserving locality: ensuring that interactions between two users in Boston, for example, do not usually – and preferably never – incur latency or bandwidth costs of communication with a data center in Japan. Many scalable systems such as Amazon’s EC2 [3] attempt to improve locality by focusing client load on “nearby” data centers and minimize wide-area interactions [34]. These mechanisms are usually ad hoc and specific to particular algorithms, however, and rarely offer quantifiable guarantees on “how well” they preserve locality, or bounds on the worst-case costs they might incur in doing so.

In pursuit of a more general, “first-principles” approach to preserving locality in distributed systems, we propose Crux, a scheme to “transform” a scalable but non-locality-preserving distributed algorithm \mathcal{A} into a similarly scalable *and* locality-preserving algorithm $\bar{\mathcal{A}}$. Crux’s transform applies to distributed algorithms satisfying a few assumptions: (a) that the underlying algorithm \mathcal{A} already scales efficiently to large clusters (e.g.,

with polylogarithmic overheads); (b) that whatever *requests* algorithm \mathcal{A} supports (e.g., GET and PUT for a DHT) can be replicated and parallelized across multiple instances of the underlying algorithm; (c) that the worst-case delays \mathcal{A} incurs servicing these requests are proportional (within polylogarithmic factors) to the *delay diameter* of the cluster on which a particular instance of \mathcal{A} runs; and (d) that \mathcal{A} effectively balances load across the cluster supporting that instance.

Given an underlying algorithm \mathcal{A} meeting these assumptions, Crux builds a collection of distributed instances of \mathcal{A} , and distributes requests across these instances, using structuring techniques inspired by landmark [30] and compact routing schemes [28, 29]. When any two users u and v interact via requests they submit to the transformed system $\bar{\mathcal{A}}$, Crux ensures that the worst-case delays incurred in servicing these requests are proportional to the network delay between u and v . This property guarantees that whenever u and v are topologically close (e.g., both in Boston), their interactions incur proportionally low delays, and never require round-trips to Japan for example – as might be the case using one instance of \mathcal{A} distributed worldwide.

At the heart of Crux is a multi-level exponential probability distribution inspired by skip-lists [24] and compact graph processing schemes [28, 29]. As in those schemes, Crux treats all participating machines as nodes in a connected graph and probabilistically assigns a *level* to each node such that higher levels contain fewer nodes. Then, Crux builds subgraphs of shortest paths trees around nodes ensuring that any two nodes always share at least one subgraph (i.e., a shortest paths tree). As a result, higher-leveled nodes act as universal *landmarks* around which distant nodes can always interact, while lower-leveled nodes act as local landmarks, incurring latency proportional to the inter-node distance.

We have implemented Crux-transformed versions of three well-known distributed systems: (a) the **memcached** caching service [9], (b) the **bamboo** distributed hash table (DHT) [4], and (c) a **redis** publish/subscribe service [25]. Experiments with roughly 100 participating machines on PlanetLab [7] demonstrate orders of magnitude improvement in interaction latency for nearby nodes, reducing some interaction latencies from roughly 900ms to 1ms. Observed latencies are typically much

lower than the upper bounds Crux guarantees, and often close to the experimentally observed optimum.

The price of Crux’s locality-preservation guarantee is that each node must participate in multiple instances of the underlying algorithm – typically 15 and up to 34 instances in our experiments, though these costs depend on tunable parameters. Nodes must also handle correspondingly higher (though evenly balanced) loads, as Crux replicates all requests to multiple instances. While these costs are significant, we anticipate they will be surmountable in settings where the benefits of preserving locality are important enough to justify their costs. For example, online gaming systems may wish to ensure low response times when geographic neighbors play each other; in earthquake warning or other disaster response systems it can be crucial for nearby nodes to receive notice of an event quickly; and modern Internet and cloud infrastructures incorporate control-plane systems in which request volume is low but response latency is crucial for localized events such as outage notifications.

This paper makes four main contributions: 1) a general framework for transforming a scalable distributed system into a locality-preserving scalable distributed system, 2) a practical application of compact graph processing schemes to limit interaction latency to a constant factor of underlying latency, 3) a prototype implementation demonstrating the feasibility of the approach, and 4) several globally distributed experimental deployments and an analysis of their performance characteristics.

Section 2 presents our assumptions about \mathcal{A} and the underpinnings of Crux’s distributed landmark hierarchy, and Section 3 details how we used Crux to build three prototype locality-preserving systems. Section 4 examines our proof-of-concept prototype implementation and Section 5 evaluates it experimentally. Section 6 discusses related work and Section 7 concludes.

2 The Crux of Locality Preservation

We now present the Crux architecture in three steps. First, we lay out the goals and key assumptions Crux makes about the underlying non-locality-preserving algorithm \mathcal{A} . Second, we present a scheme to transform \mathcal{A} into an intermediate, *landmark-centric* algorithm $\bar{\mathcal{A}}$, which “preserves locality” in a restrictive setting where all distances are measured to, from, or via some *landmark* or reference point. Finally, by treating *all* nodes as landmarks of varying levels chosen from a suitable random distribution [28, 29], and building varying-size instances of $\bar{\mathcal{A}}$ around these landmarks, we transform $\bar{\mathcal{A}}$ into the final algorithm $\bar{\bar{\mathcal{A}}}$, which guarantees locality among all interacting pairs of nodes while keeping overheads polylogarithmic at every node.

2.1 Goal and Success Metric

Consider a simple publish/subscribe service as implemented by **redis** [25], for example, distributed across multiple nodes relaying information via named *channels*. Suppose this pub/sub cluster spans servers around the world, and that a user Bob in Boston publishes an event on a channel that Charlie in New York has subscribed to. We say that Bob and Charlie *interact* via this **redis** pub/sub channel, even if Bob and Charlie never communicate directly at network level. Because the **redis** server handling this channel is chosen randomly by hashing the channel name, this server is as likely to be in Europe or Asia as on the US East Coast, potentially incurring long interaction latencies despite the proximity of Bob and Charlie to each other. To *preserve locality* we prefer if Bob and Charlie could interact via a nearby **redis** server, so that their interactions via this pub/sub service never require a round-trip to Europe or Asia.

To increase locality, realistic service deployments often create separate service clusters for each data center or geographic region and then “assign” users to clusters so that users in the same region interact locally. Typically there is no way to measure *how* well they preserve locality, however, let alone any quantifiable locality *guarantees*. Region selections are typically ad hoc or arbitrary, and can have surprising side-effects on some “unlucky” users. For example, between two users in the central US who are very close to *each other*, but just across a service’s “border” between US East and US West regions, interactions may require round-trip communications involving *both* east-coast and west-coast servers and exhibit “worst of both worlds” latencies.

The goal of Crux, therefore, is to develop a general, “first principles” approach to building distributed services that preserve locality. To measure success, we use as our “yardstick” the network delay with which two users – e.g., Bob and Charlie in the above example – would observe *if* they were communicating directly via the underlying network (e.g., with peer-to-peer TCP connections). For *any* distributed service of arbitrary internal complexity and satisfying certain requirements detailed next, we seek a systematic way to deploy this service over a global network, while ensuring that two users’ *interaction latency* through the service remains proportional (within polylogarithmic factors) to their point-to-point communication latency over the underlying network directly. One of this problem’s key technical challenges is to preserve locality for *all* interacting pairs of users, even in corner cases, such as when users are “just across a region border.”

2.2 The Underlying Algorithm \mathcal{A}

Crux works with any distributed algorithm \mathcal{A} satisfying certain assumptions. We define these assumptions infor-

mally because they represent properties whose precise formal definition – if one exists – tends to vary from one distributed algorithm to the next.

Network: We assume the final locality-preserving algorithm $\bar{\mathcal{A}}$ is expected to run on some distributed network represented by a graph (V, E) consisting of nodes $v \in V$, and edges $e \in E$. Nodes typically represent physical hosts, such as peers in a P2P system or servers composing a distributed cloud service. Edges may represent physical links, TCP connections, tunnels, or any other logical communication path. We assume the network is connected but otherwise of arbitrary topology, and that every node may communicate with any other (i.e., multi-hop routing is already in place and reliable).

Distance: We assume every pair of nodes $u, v \in V$ have a *distance*, which we denote \overline{uv} . Crux can use any distance metric, but for clarity we assume distances represent round-trip communication delay between u and v over the underlying network. Two nodes communicating over multi-hop routes need not necessarily use shortest-path routes, but we do assume that the triangle inequality holds: i.e., $\overline{uv} \leq \overline{uw} + \overline{wv}$ for any nodes u, v, w . This assumption is not always true on the Internet, unfortunately [17], but triangle inequality violations should only break Crux’s locality preservation guarantees in proportion to the violation’s severity.

Requests: We assume the underlying algorithm \mathcal{A} is service-oriented, as per the “Bezos Mantra” [32]: that \mathcal{A} exists to satisfy RPC-style requests. For simplicity we assume requests originate internally from nodes in V , as in a P2P system in which all clients are also servers, or in a cloud service whose clients reside in the same infrastructure. Requests might alternately originate from a separate set of external clients, but we assume such requests will arrive *via* some set of front-end or proxies that are part of V . Regardless of \mathcal{A} ’s semantics or the specific types of requests it supports, we assume these requests may be classified into *read*, *write*, or *read/write* categories. Read requests retrieve information about state managed by \mathcal{A} , write requests insert new information or modify existing state, and read/write requests may both read existing state and produce new or modified state.

Interaction: We assume the end goal of \mathcal{A} is to enable groups of two or more users, represented by network nodes, to *interact* indirectly through requests they submit to some shared instance of \mathcal{A} . We primarily consider interactions between *pairs* of nodes u and v , treating a larger interacting group as an aggregate of all pairs of group members. If \mathcal{A} is a DHT, for example, nodes u and v might interact when u PUTs a key/value pair that v subsequently GETs. If \mathcal{A} is a distributed publish/subscribe service, u and v interact when u publishes an event on a

channel that v subscribes to. In this way, interactions are logical, not physical; the DHT example might satisfy v ’s GET for u ’s data via some other node w .

Of course, instead of interacting through \mathcal{A} , u might in principle simply send a message directly to v on the underlying network, with round-trip delay \overline{uv} as discussed above. The purpose of higher-level distributed services as abstracted by \mathcal{A} , however, is to support more complex forms of interaction not easily satisfied merely by point-to-point messaging – such as when u and v are not directly aware of each other but only know a common key or publish/subscribe channel. Nonetheless, we will adopt the baseline communication delay between u and v as a *measuring stick*, against which we compare the latency of the final locality-preserving algorithm $\bar{\mathcal{A}}$.

Instantiation: We assume that we can create and manage multiple independent *instances* of the underlying distributed algorithm \mathcal{A} , and can deploy each instance i on any subset of nodes $V_i \subseteq V$. We further assume each physical node may participate in multiple instances of \mathcal{A} simultaneously: two instances i, j of \mathcal{A} may be distributed across overlapping sets of nodes V_i, V_j . This is true in practice provided the relevant software is configurable to use disjoint node-local resources for each instance: e.g., separate configuration and state directories and disjoint TCP/UDP ports. Crux does not require these instances of \mathcal{A} to coordinate or be “aware of” each other. While such coordination could enable optimization opportunities or stronger consistency models, in our basic model Crux manages all coordination across instances.

Scalability: We assume that \mathcal{A} already scales efficiently, at least up to the size of the target network $N = |V|$. We do not attempt to improve \mathcal{A} ’s scalability, but merely preserve its scalability while introducing locality preservation properties. We intentionally leave the precise definition of “scalability” informal since in practical systems it reflects how gracefully many different measurable properties or overheads grow with network size: e.g., the average or worst-case growth rates of each node’s CPU load, memory and/or disk storage overhead, network bandwidth utilization, request processing delays, etc. However, we assume that for whatever specific overheads are of interest, \mathcal{A} is “scalable” if those overheads grow polylogarithmically with network size. Thus, Crux’s scalability goal is to ensure that the corresponding overheads in the final $\bar{\mathcal{A}}$ similarly grow polylogarithmically in network size.

Performance: We assume that request processing time is the primary performance metric of interest, and that \mathcal{A} ’s processing time for any request grows proportionally to the maximum communication delay among the nodes across which \mathcal{A} is distributed. Specifically, Crux assumes that for any given request r submitted to an \mathcal{A}

instance i deployed on node set V_i , \mathcal{A} can process r in time $T = \tilde{O}(D_i)$, where D_i is the *diameter* of the sub-graph V_i : the worst-case round-trip time between any two nodes $u, v \in V_i$. This assumption that \mathcal{A} 's processing time is proportional to its *longest-path* communication delay follows from the assumption that \mathcal{A} is *not* locality-preserving. As with scalability metrics, Crux is ambivalent about whether the latency metrics of interest represent average-case, worst-case, or 90-percentile cases: for any latency metric for which \mathcal{A} is “network-bottlenecked” and hence related to delay diameter as discussed above, Crux ensures that the resulting $\bar{\mathcal{A}}$ preserves locality in the corresponding latency metric.

Request replication: We assume that it is both possible and compatible with \mathcal{A} 's semantics, to *replicate* any incoming request by directing copies of the same request to multiple instances of \mathcal{A} , to be chosen by Crux as detailed later. Request replication might be managed by Crux-aware clients themselves, or by one or more intermediate proxy or load-balancing nodes in the network. We assume a *read* request may be satisfied using the *first* “success” response from any of the \mathcal{A} instances to which the read request is replicated. We assume a *write* request will be processed to completion by *all* instances of \mathcal{A} to which Crux replicates the request. However, we assume the relevant state insertion or modification caused by the request will “take effect” and become available for subsequent reads within each \mathcal{A} instance as soon as the write request directed to *that instance* completes. That is, if Crux replicates a write request to both a “small” \mathcal{A} instance i on nodes V_i with low diameter D_i , and a “large” instance j on nodes V_j with high diameter $D_j \gg D_i$, both replicas of the write request will eventually complete (in time proportional to D_j), but the write will “take effect” and become available in instance i in time proportional to D_i . For any pair of nodes $u, v \in V$ interacting via the system, we assume their interaction will be successful provided that Crux consistently directs u 's and v 's requests to at least one common instance of \mathcal{A} .

Consistency: Our assumption that requests may be replicated freely and obviously to multiple instances of \mathcal{A} , implies an assumption that \mathcal{A} embodies a weak consistency model. While we expect that well-known distributed consensus techniques could be adapted to the Crux model to offer locality-preserving service with stronger consistency, we defer exploration of this challenge to future work.

2.3 Landmark-Centric Locality

Crux's first step is to transform the initial non-locality-preserving algorithm \mathcal{A} into an algorithm $\bar{\mathcal{A}}$ that provides *landmark-centric locality*: that is, locality when all distances are measured to, from, or through a well-defined




Figure 1: Dividing a network into concentric rings around a landmark L to achieve *landmark locality*.

reference point or *landmark*. We make no claim that this transform is novel in itself, as many existing distributed algorithms embody similar ideas [30, 31]; we merely utilize it as one step toward our goal of building a general scheme that enables locality-preserving systems.

Rings: This first transformation is conceptually simple: given an arbitrary node L chosen to serve as the landmark, Crux divides all network nodes $v \in V$ into concentric *rings* or *balls* around L , as illustrated in Figure 1, assigning nodes to instances of \mathcal{A} according to these rings. Successive rings increase exponentially in radius (distance from L), from some minimum radius r_{min} to a maximum radius r_{max} large enough to encompass the entire network. The minimum radius r_{min} represents the smallest “distance granularity” at which Crux preserves locality: Crux treats any node u closer than r_{min} to L as if it were at distance r_{min} . We define the ratio $R = r_{max}/r_{min}$ to be the network’s *radius spread*. We will assume R is a power of 2, so the maximum number of rings is $\log_2(R)$. For simplicity we subsequently assume network distances are normalized so that $r_{min} = 1$, and hence $R = r_{max}$.

Instantiating $\bar{\mathcal{A}}$: Given these distance rings around L , Crux creates a separate instance of \mathcal{A} for each non-empty ring. At this point we have a design choice between *exclusive* or *inclusive* rings, yielding tradeoffs discussed below. With *exclusive rings*, each network node belongs to *exactly* one ring, and hence participates in exactly one instance of \mathcal{A} . With *inclusive rings*, in contrast, outer rings “build on” inner rings and include all members of all smaller rings, so that nodes in the innermost ring (e.g., L and B in Figure 1) participate in up to $\log_2(R)$ separate instances of \mathcal{A} . The choice of inclusive rings obviously increases the resource costs each node incurs – by a factor of $\log_2(R)$ for all “per-instance” costs such as memory, setup and maintenance bandwidth, etc. We expect inclusive rings to exhibit more predictable load

and capacity scaling properties as discussed later in Section 2.5, however, so we will henceforth assume inclusive rings except when otherwise stated.

Replicating requests: Whenever any node u initiates a service request to $\bar{\mathcal{A}}$ (or whenever a request from an external client arrives via some ingress node u), Crux replicates and directs copies of the request to the instances of \mathcal{A} representing u 's own ring and all larger rings. Consider for example the case of a simple distributed key/value store based on consistent hashing, such as **memcached** [9]. If node D in Figure 1 issues a *write* request for some key k and value v , then Crux issues copies of this request to the two (logically independent) **memcached** instances representing rings 1 and 2. In this case, key k hashes to node E in ring 1 and to node F in ring 2, so this replicated *write* leaves copies of (k, v) on nodes E and F . If node B in ring 0 subsequently issues a *read* request for key k , Crux replicates this *read* to all three rings: in this case to node C in ring 0, which fails to find (k, v) , and in parallel to nodes E and F , which do (eventually) succeed in finding the item. Since copies of all requests are directed to the largest ring, any pair of interacting nodes will successfully interact via that outer ring if not (more quickly) in some smaller interior ring.

Delay Properties of $\bar{\mathcal{A}}$ The key property $\bar{\mathcal{A}}$ provides is *landmark-centric locality*. Specifically, for any two interacting nodes u and v at distances \overline{uL} and \overline{vL} from L , respectively, $\bar{\mathcal{A}}$ can process the interaction's requests in time proportional to $\max(\overline{uL}, \overline{vL})$. By the assumptions in Section 2.2, each instance i of the underlying algorithm \mathcal{A} can process a request in time $\tilde{O}(D_i)$, where D_i is the diameter of the set of nodes V_i over which instance i is distributed. Assuming inclusive rings, for any nodes u, v there is some innermost ring i containing both u and v . The diameter of ring i is 2×2^i , so the ring i instance of \mathcal{A} can process requests in time $\tilde{O}(2^i)$. Assuming $r_{\min} = 1 \leq \max(\overline{uL}, \overline{vL})$, either u or v must be a distance of at least 2^{i-1} from L , otherwise there would be a ring smaller than i containing both u and v . Hence, $2^{i-1} \leq \max(\overline{uL}, \overline{vL})$, so $\bar{\mathcal{A}}$ can process the requests required for u and v to interact in time $\tilde{O}(\max(\overline{uL}, \overline{vL}))$.

Consider the example in Figure 1, supposing \mathcal{A} is **memcached**. While node D 's overall *write*(k, v) may take a long time to insert (k, v) in rings farther out, by the assumptions stated in Section 2.2 (which are valid at least for **memcached**), the copy of this *write* directed at ring 1 will hash to *some* node of distance no more than 2^1 from L , node E in this case. That (k, v) pair then becomes available for reading at node E in time proportional to ring 1's radius. The subsequent *read*(k) by node B will fail to find k in ring 0, but its *read* directed at ring 1 will hash to node E and find the (k, v) pair inserted by node D in time proportional to ring 1's radius.




Figure 2: Dividing a network into overlapping *clusters* around multiple landmarks to achieve all-pairs locality preservation. Each cluster's size is constrained by *higher-level* landmarks; top-level clusters (B_2) cover the whole network. Nodes (e.g., A) replicate requests to their own and outer rings for all landmarks in their bunch.

If the request requires only a single communication round-trip, as may be the case in a simple key/value store like **memcached**, then the total response time is bounded by the diameter of the ring i through which two nodes interact. If \mathcal{A} requires more complex communication, requiring a polylogarithmic number of communication steps within ring i , for example – such as the $O(\log N)$ steps in a DHT such as Chord [27] – then $\bar{\mathcal{A}}$'s ring i still bounds the network delay induced by each of those communication hops *individually* to $\tilde{O}(2^i)$, yielding an overall processing delay of $\tilde{O}(2^i) = \tilde{O}(\max(\overline{uL}, \overline{vL}))$.

2.4 All-Pairs Locality Preservation

We now present a transform yielding an algorithm $\bar{\mathcal{A}}$ that preserves locality among all interacting pairs of nodes. The basic idea is to instantiate the above landmark-centric locality scheme many times, treating *all* nodes as landmarks distributed as in a compact routing scheme [28, 29]. This process creates overlapping sets of ring-structured instances of \mathcal{A} around *every* node, rather than just around a single designated landmark. This construction guarantees that for every interacting pair of nodes u, v , both nodes will be able to interact through *some* ring-instance that is both “small enough” and “close enough” to both u and v to ensure that operations complete in time proportional to \overline{uv} . Despite producing a large (total) number of \mathcal{A} instances network-wide, most of them will contain only a few participants and, thus, each node will need to participate in a relatively small (polylogarithmic) number of instances of \mathcal{A} .

The Landmark Hierarchy: We first assign each node $v \in V$ a *landmark level* l_v from an exponential distribution with a *base* parameter B . We initially assign all nodes level 0, then for each successive level i , we choose each level i node uniformly at random with probability $1/B$ to be “upgraded” to level $i + 1$. This process

stops when we produce the first empty level k , so with high probability $k \approx \log_B(N)$. Figure 2 illustrates an example landmark hierarchy, with a single level 2 landmark B_2 , two level 1 landmarks C_1 and D_1 , and all other nodes (implicitly) being level 0 landmarks.

Bunches: For every node u in the network, we compute a *bunch* B_u representing the set of surrounding landmarks u will be aware of, as in Thorup/Zwick compact routing [28]. To produce u 's bunch, we search outward from u in the network, conceptually encountering every other node v in ascending order of its distance from u , as in a traversal using Dijkstra's single-source shortest-path algorithm with u as the source. For each node v we encounter in this outward search, we include v in u 's bunch if v 's landmark level l_v is *no smaller* than that of any node we have encountered so far. Stated another way, if we assume u 's landmark level is 0 then in the outward search we include in u 's bunch all the level 0 nodes until (and only until) we encounter the first node at level 1 or higher. Then we include all level 1 nodes – while ignoring all subsequent level 0 nodes – until we encounter the first node at level 2 or higher, and so on until we encounter all nodes. In Figure 2, for example, node A includes landmarks C_1 and B_2 in its bunch, but *not* D_1 , which is farther from A than B_2 is.

Due to the landmark level assignment method, we expect to accept about B nodes into u 's bunch at each level i , before encountering the first level $i+1$ node and subsequently accepting no more level i nodes. Since there are about $\log_B(N)$ levels total, with high probability each node's bunch will be of size $|B_u| \approx B \log_B(N)$, using basic Chernoff bounds. This is the key property that enables us to bound the number of instances of \mathcal{A} that node u will ultimately need to be aware of or participate in.

Clusters and Instantiating $\hat{\mathcal{A}}$: For every node v having landmark level l_v , we define v 's *cluster* as the set of nodes $\{u_1, u_2, \dots\}$ having v in their bunch: i.e., the set of nodes around v that are close enough to “know about” v as a landmark. Node v is trivially a member of its own cluster. For each node v , we apply the landmark-centric locality scheme from Section 2.3 to all nodes in v 's cluster, using v as the landmark L around which $\hat{\mathcal{A}}$ forms its up to $\log_2(R)$ ring-structured instances of \mathcal{A} .

For each node u in v 's cluster, u will be a member of at least one of v 's ring-instances of \mathcal{A} , and potentially more than one in the case of inclusive rings. We provide each such node u with the network contact information (e.g., host names and ports) necessary for u to submit requests to the instance of \mathcal{A} representing u 's own ring, and to the instances of \mathcal{A} in all larger-radius rings around v . In Figure 2, for example, B_2 's top-level cluster encompasses the entire network, while the clusters of lower-level land-

marks such as C_1 include *only* nodes closer to C_1 than to B_2 (e.g., A).

Per-Node Participation Costs: Since every node $v \in V$ is a landmark at *some* level, we ultimately deploy $N = |V|$ total instances of $\hat{\mathcal{A}}$, or up to $N \log_2(R)$ total instances of the original algorithm \mathcal{A} throughout the network. Each node u 's bunch contains at most $\approx B \log_B(N)$ landmarks $\{v_1, v_2, \dots\}$, and for each such landmark $v \in B_u$, u maintains contact information for at most $\log_2(R)$ instances of \mathcal{A} forming rings around landmark v . Thus, each node u is aware of – and participates in – at most $\approx B \log_B(N) \log_2(R)$ instances of \mathcal{A} total. This property ultimately ensures that the overheads $\hat{\mathcal{A}}$ imposes at each node on a per- \mathcal{A} -instance basis are asymptotically $\tilde{O}(1)$ or polylogarithmic in both the network's total number of nodes N and radius spread R .

Handling Service Requests: When any node u introduces a service request, u replicates and forwards simultaneous copies of the request to all the appropriate instances of \mathcal{A} , as defined by $\hat{\mathcal{A}}$, around *each* of the landmarks in u 's bunch. Since u 's bunch contains at most $\approx B \log_B(N)$ landmarks with at most $\approx \log_2(R)$ rings each, u needs to forward at most $\approx B \log_B(N) \log_2(R)$ copies of the request to distinct instances of \mathcal{A} . As usual, we assume these requests may be handled in parallel. In Figure 2, node A replicates its *read*(k) request to the \mathcal{A} instances representing rings 1 and 2 around landmark C_1 (but not to C_1 's inner ring 0). A also replicates its request to rings 2 and 3 of top-level landmark B_2 , but not to B_2 's inner rings.

In practice, a straightforward and likely desirable optimization is to “pace” requests by submitting them to nearby landmarks and rings first, submitting requests to more distant landmarks and rings only if the nearby requests fail or time-out. For example, “expanding-ring” search methods – commonly used in ad-hoc routing [22] and peer-to-peer search [18] – can preserve Crux's locality guarantees within a constant factor provided that the search radius increases exponentially. For conceptual simplicity, however, we assume that u launches all replicated copies of each request simultaneously.

Interaction Locality: Using the same reasoning as in Thorup/Zwick's stretch $2k - 1$ routing scheme [29], the landmark assignment and bunch construction above guarantee that for *any* two nodes u, v in the network, there will be *some* landmark L present in both u 's and v 's bunches such that $\overline{uL} + \overline{vL} \leq (2k - 1)\overline{uv}$. This key locality guarantee relies only on the triangle inequality and otherwise makes no assumptions about network topology, and hence holds on any graph. Further, this property represents a worst-case bound on *stretch*, or the ratio between the distance of the *indirect* route via L (i.e., $\overline{uL} + \overline{vL}$), and the distance of the *direct* route \overline{uv} . Practi-

cal experience with this and similar schemes has shown that in typical networks, stretch is usually much smaller and often close to 1 (optimal) [14, 15].

Thus, $\bar{\mathcal{A}}$'s landmark hierarchy ensures that u and v can always interact via *some* common landmark, L , using one of its $\bar{\mathcal{A}}$ ring instances. Further, u 's and v 's landmark-centric distances \overline{uL} and \overline{vL} are no more than $O(k) = O(\log N) = \tilde{O}(1)$ longer than the direct point-to-point communication distance between u and v . $\bar{\mathcal{A}}$ in turn guarantees that the instance of \mathcal{A} deployed in the appropriate ring around L will be able to process such requests in time proportional to $\max(\overline{uL}, \overline{vL})$. Thus, interactions between u and v via the selected ring ultimately take time $\tilde{O}(\overline{uL})$, yielding the desired locality property.

2.5 Load and Capacity Considerations

As noted earlier, we can choose either inclusive or exclusive rings in the intermediate landmark-centric algorithm $\bar{\mathcal{A}}$. Exclusive rings have the advantage that each node in an $\bar{\mathcal{A}}$ instance participates in *exactly one* ring around the landmark L , and hence in only one \mathcal{A} instance per landmark. Inclusive rings require nodes in the inner rings to participate in up to $\log_2(R)$ rings around L , and hence in multiple distinct instances of \mathcal{A} per landmark.

Potentially compensating for this increased participation cost, however, the inclusive ring design has the desirable property that every node u that may incur system load by introducing requests, replicates those requests *only* to instances of \mathcal{A} in which u itself participates. If u 's bunch includes landmark L and u is in L 's ring i , for correctness u must replicate its requests to all of L 's rings numbered i or higher as discussed above.

With inclusive rings, the request-processing workload imposed on any particular instance of \mathcal{A} comes *only* from members of that same instance of \mathcal{A} , and hence $\bar{\mathcal{A}}$ preserves \mathcal{A} 's capacity-scaling properties. Assume that each node introduces requests at some maximum rate r – either “naturally” or by explicit rate-limiting – and that an \mathcal{A} instance distributed over n nodes scales gracefully to handle requests at rates up to $\approx nr$. (Any relevant load metric may be used; we assume request rate only for simplicity.) With inclusive rings, each \mathcal{A} instance within an $\bar{\mathcal{A}}$ network, of some size n , services *only* requests originating from within its n -node membership, ensuring that it need handle at most a $\approx nr$ total request rate – which is the case if \mathcal{A} 's capacity scales across those n nodes.

With exclusive rings, members of inner rings around some landmark L must still submit requests to outer rings they are not members of. Thus, some highly populous inner ring might overload a much sparser outer ring instance with more requests than the outer instance of \mathcal{A} can collectively handle. With inclusive rings such capacity imbalance does not arise because all the members of the populous inner ring are *also* members of the larger

outer rings and, by virtue of their membership, increase the outer rings' capacity to handle the inner rings' imposed load.

2.6 Asymmetric Request Replication

In the basic $\bar{\mathcal{A}}$ above, each node u indiscriminately replicates *all* requests to the $\bar{\mathcal{A}}$ instances around *all* landmarks L in u 's bunch. Suppose however: (a) \mathcal{A} 's requests are divided into two classes – we will assume *reads* and *writes* for simplicity; (b) one class is more expensive to replicate, e.g., *writes*, as they consume storage; and (c) for correctness \mathcal{A} requires only that any node's *read* requests will “meet” any other interacting node's *write* requests in some common \mathcal{A} instance. In particular, we now assume there is no inherent need for two *write* requests to “meet” in an \mathcal{A} instance, only that *reads* will be able to “find” relevant *writes*. In this case, it is sufficient for node u to replicate its *write* requests only to the landmark L_i in u 's bunch *closest to* u at each level i .

With this change, *write* requests are replicated at most $\approx \log_B(N) \log_2(R)$ times, a constant factor B lower than read request replication. Crux still guarantees that *read* requests will meet *write* requests at some \mathcal{A} instance, and still preserves locality, but in this case the upper bound in worst-case stretch may be higher by a factor of two. In particular, the formal reasoning underlying this approach corresponds to the “no handshaking” variant of Thorup/Zwick's algorithm [29] (Appendix A), which bounds stretch only to $4k - 3$ rather than $2k - 1$.

3 Applying Crux in Realistic Systems

We now explore ways to apply the Crux transform to practical distributed algorithms. We approach the task of implementing Crux as an initial *instance assignment* process, followed by a dynamic *request replication* process during system operation. We summarize how we applied Crux to build three example locality-preserving systems based on **memcached**, **bamboo**, and **redis**.

3.1 Instance Assignment

We treat instance assignment as a centralized pre-processing step, which might be deployed as a periodic “control-plane” activity in today's data center or software-defined networks [21], for example. While we expect it should be possible to build fully distributed, “peer-to-peer” Crux systems requiring no central processing, we defer this challenge to future work.

Crux instance assignment takes as input a suitable *network map*, represented as a matrix of round-trip latencies, or some other suitable metric, between all pairs of nodes u and v . In our ≈ 100 -node prototype deployments, a centralized server simply runs a script on each node u measuring u 's minimum ping latency to all other

nodes v . We expect mature deployments to use more efficient and scalable network mapping techniques [19], but do not address this problem here.

Crux processes this network map first by assigning each node u a landmark level, then computing u 's bunch B_u and cluster C_u according to u 's level and distance relation to other nodes. As described in Section 2.4, Crux uses the computed bunches to identify which landmarks' ring instances each node will join.

For every node u in node v 's cluster C_v (including v itself), Crux assigns u to a single ring (with exclusive rings) or a set of consecutive rings (with inclusive rings) around v . Across all nodes, each non-empty ring constructed in this fashion corresponds to a unique instance of \mathcal{A} in which the assigned nodes participate. Crux then transmits this instance information to all nodes, at which point they launch or join their \mathcal{A} instances.

3.2 Request Replication in $\bar{\mathcal{A}}$

After computing and launching the appropriate \mathcal{A} instances, Crux must subsequently interpose on all service requests introduced by (or via) member nodes, and replicate those requests to the appropriate set of \mathcal{A} instances.

In our prototype, the central controller deploys on each node u not just the instance configuration information computed as described above, but also a *trigger* script specific to the particular algorithm \mathcal{A} . This trigger script emulates the interface a client normally uses to interact with an instance of \mathcal{A} , transparently interposing on requests submitted by clients and replicating these requests in parallel to all the \mathcal{A} instances u participates in. This interposing trigger script thus makes the full $\bar{\mathcal{A}}$ deployment appear to clients as a single instance of \mathcal{A} .

Although many aspects of Crux are automatic and generic to any algorithm \mathcal{A} , our current approach does require manual construction of this algorithm-specific trigger script. As noted earlier in Section 2.2, our trigger scripts also effectively assume \mathcal{A} requires and offers only weak consistency, so that the trigger script's parallel replication of all requests create no problematic side-effects. For operations returning a value – (e.g., *read*) – the trigger script returns to the user the first value that it successfully retrieves.

3.3 Example Applications of Crux

Crux intentionally abstracts the notion of an underlying scalable distributed algorithm \mathcal{A} , so as to transform a general class of distributed systems into a locality-preserving counterpart $\bar{\mathcal{A}}$. This section discusses our experience transforming three widely-used distributed services. We deployed all three services with no modifications to the source code of the underlying system.

Memcached: We used **memcached** [9] as an example of a distributed key/value caching service, where a col-

lection of servers listens for and processes *Put* and *Get* requests from clients. A *Put* stores a key/value pair at a server, while a *Get* fetches a value by its associated key. Clients use a consistent hash on keys, modulo the number of servers in a sorted list of servers, to select the server responsible for a particular key. Servers do not communicate or coordinate with each other at all, which is feasible for **memcached** due to its best-effort semantics, which make no consistency guarantees.

We chose to deploy **memcached** using Crux because of its general popularity with high-traffic sites [9] and its simple interaction model. In our deployment, we map each \mathcal{A} ring instance in a node's cluster to a collection of servers acting as a single **memcached** instance. Each node's trigger script simply maintains a separate set of servers for each of its instances and then uses a consistent hash to select the server to handle each request.

Bamboo: We next built a locality-preserving deployment of **bamboo** [4], a distributed hash table (DHT) descended from Pastry [26], relying on the same Crux instance membership logic as in **memcached**. We chose a DHT example because of their popularity in storage systems, and because DHTs represent more complex “multi-hop” distributed protocols requiring multiple internal communication rounds to handle each client request. Instantiating a **bamboo** instance requires the servers to coordinate and build a distributed structure, but clients need only know of *one* server in the DHT instance in order to issue requests. A DHT like **bamboo** may in principle scale more effectively than **memcached** to services distributed across a large numbers of nodes, since each client need not obtain – or keep up-to-date – a “master list” of all servers comprising the DHT instance.

Redis: Lastly, we built a locality-preserving publish/subscribe service built on the pub/sub functionality in **redis** [25]. In this pub/sub model, clients may *subscribe* to named *channels*, and henceforth receive copies of messages other clients *publish* to these channels. As in the **memcached** example, clients use consistent hashing on the channel name to select the **redis** server (from a sorted server list) responsible for a given channel within an \mathcal{A} instance. When a client subscribes to a particular channel, it maintains a persistent TCP connection to the selected **redis** server so that it can receive published messages without polling.

3.4 Potential Optimizations

We now briefly point out potential further optimizations, though our prototype does not yet implement them.

Network mapping: Rather than mapping the underlying network by ping-testing all node pairs, nodes could maintain and refer on-demand to a suitable cache of net-

	Shell	Python	C/C++
Graph Production	51	165	
Graph Processing			1,041
Deployment		2,486	

Table 1: Crux prototype code size. Distributed service binaries were used as-is, without modification.

work map and distance information, such as that maintained by a service like iPlane [19].

Landmark selection: Depending on the costs and behavior of \mathcal{A} instances, some nodes may experience a higher load than others due to being a member of more instances. Crux might be able to alleviate the load on weak or overloaded nodes by changing either the level assignment mechanism, or the “capacity share” each node is responsible for within a given \mathcal{A} instance.

Pacing requests: Perhaps the most beneficial optimization may be for the trigger script to query smaller, closer \mathcal{A} instances first and larger instances only if these first requests fail or time-out, reducing the total number of replicated requests the system must handle especially if a considerable fraction of traffic is localized.

4 Prototype Implementation

Our current prototype comprises three components written in a mix of shell scripts, Python and C/C++. As written, the code is generally unoptimized, and we have no expectation that it represents the “best” way to implement or deploy the Crux architecture. Table 1 summarizes the language and code complexity of each component in our current implementation in lines of code [8].

The graph production component, written in Python and shell scripts, takes as input a list of hosts on which to deploy. A central server produces a graph containing vertices (hosts) and edges (network propagation delay between hosts). This process has running time on the order of the size of the graph and is heavily influenced by the network delay between hosts. For the roughly 100-node deployment used in the experimental evaluation, generating the graph took less than one hour for a maximum observed inter-node latency of ≈ 500 ms.

The graph processing component, written in C and C++, reads the input graph, assigns a level to each node and produces the bunches and clusters forming the distributed cover tree. Graph processing has running time on the order of the number nodes and is heavily influenced by the maximum landmark level constant, k , where a smaller k increases the processing time. For 100 nodes and $k = 2$, processing takes roughly 1 second, while processing a graph of 350,000 nodes with k as small as 4 completes in roughly 30 seconds.

The deployment component, written in Python, first translates the bunches and clusters from the graph pro-




Figure 3: Average bunch size per node for a 96-node graph as the level constant, k , increases.

cessing stage into \mathcal{A} instance membership lists. Next, it triggers each participating node to download a zipped archive containing the instance membership lists, the trigger script and executables of the underlying \mathcal{A} implementation. A local script launches or joins any of the node’s member instances, after which the trigger script begins listening for and servicing operations. Our current prototype supports only batch operations for testing and does not include a “drop-in” replacement user interface for generic client use.

5 Experimental Evaluation

This section evaluates our Crux prototype, first by examining the effect of the level constant k on per-node storage costs, and second, in large-scale distributed deployments on PlanetLab [7]. All experiments use a network graph consisting of 96 PlanetLab nodes. The live experiments use inclusive rings except where otherwise noted, and measure the end-to-end latency of interactions between pairs of nodes.

5.1 Compact Graph Properties

Recall that Crux’s graph processing component collects landmarks into each node’s bunch in order of increasing distance from the node and landmark level. The maximum level constant k plays a central role in determining the expected per node bunch size and, thus, influences the overhead of joining multiple instances of \mathcal{A} .

A smaller k creates fewer landmark levels, causing a node to encounter more nodes at its own level – adding them to its bunch – before reaching a higher-leveled node. Conversely, a larger k increases the chance that a node encounters a higher-leveled node, reducing the growth of the node’s bunch at each level. Of course, k also affects the worst-case “stretch” guarantee, and hence how tightly Crux preserves locality. A larger k increases worst-case interaction between two nodes u and v to a larger multiple of their pairwise network distance \overline{uv} .

Figure 3 demonstrates the effect that the maximum level constant k has on nodes’ average bunch size. The




Figure 4: Average \mathcal{A} instance memberships per node for a 96-node graph as the level constant, k , increases.

red line represents the expected bunch size for any node, calculated as: $kn^{1/k}$, while the green line shows the actual average bunch size per node. A k value of 1 (not shown) trivially allows for no stretch at all and requires the full $N \times N$ all-pairs shortest path calculation (every node is in every other node's bunch). When $k = 2$, all nodes are assigned level 0 or 1 and have an expected versus actual bunch size of 20 and 17, respectively. Using $k = 5$ produces the smallest actual bunch sizes for this network, at around 7, with the comparatively weaker guarantee that inter-node interaction latency could be as high as $2 \times 7 - 1 = 13 \times$ the actual pairwise latency. (Increasing k above 5 actually increases the expected bunch size because the n in $kn^{1/k}$ remains constant.)

We next consider the expected versus actual number of \mathcal{A} instances that each node must create or join based on its bunch and cluster. These benchmarks present both inclusive ring instances, where a member of one ring instance is also a member of all larger ring instances around the same landmark, and exclusive ring instances. The upper bound on actual instance memberships for inclusive rings is the expected bunch size multiplied by the maximum number of rings around a landmark: $kn^{1/k} \log_2(R)$, where R represents the radius spread as previously defined.

Figure 4 shows the expected upper bound on instance memberships versus the actual average number of instance memberships. The graph indicates that while the upper bound on instance memberships can be quite high, the average expected memberships are much lower, always less than 50 distinct \mathcal{A} instances in our experiments, regardless of exclusivity, and averaging roughly 20 instances when $k = 5$.

Figure 5 shows the cumulative distribution function of the number of instance memberships per node using a level constant $k = 5$. This second experiment uses a new random level assignment but corroborates the trends identified in Figure 4 that a majority of the 96 nodes must join less than 20 distinct \mathcal{A} instances when using $k = 5$, for a stretch factor of $2 \times 5 - 1 = 9$.




Figure 5: CDF of the number of instance memberships per node for a 96-node graph using level constant $k = 5$.

In summary, the per-instance costs of \mathcal{A} and the deployment's tolerance for additional interaction latency will determine the ideal value for k . The remaining experiments use $k = 5$.

5.2 Locality-Preserving memcached

We now evaluate the total interaction latency for a Crux deployment using **memcached** as the algorithm \mathcal{A} . We define a **memcached** interaction as a *Put* operation initiated by a node u with a random key/value pair, followed by a *Get* operation initiated by a different node v to retrieve the same key/value pair. We define the *interaction latency* as the sum of the fastest *Get* and *Put* request replicas that "meet" at a common **memcached** instance, as detailed and further discussed below.

The workload for these experiments consists of nearly 100 *Put* operations by each node using keys selected randomly from a set of all unique 2-tuples of server identifiers. For example, a node u performing a *Put* operation might randomly select the key $\langle u, v, w \rangle$, where v uniquely identifies another node (e.g., by hostname) and w is an integer identifier between 0 and 9. Node u then uses this key to insert a string value into all the \mathcal{A} instances of which it is a member (in Crux), or the single, global \mathcal{A} instance otherwise. Nodes then initiate *Get* operations, using the subset of random keys that were chosen for *Put* operations. This ensures that all *Get* requests always succeed in at least one \mathcal{A} instance.

Given the above workload, the interaction latency for a single global **memcached** instance is simply the latency of node u 's *Put* operation plus the latency of node v 's (always successful) *Get* operation. Crux of course replicates u 's *Put* operation to multiple instances, and we do not know in advance from which \mathcal{A} instance node v will successfully *Get* the key. We therefore record all *Put* and *Get* operation latencies, and afterwards identify the smallest \mathcal{A} ring instance through which u and v interacted, recording the sum of those particular *Put* and *Get* latencies as the interaction latency for Crux. Although the *Get* operation in this experiment does not occur im-




Figure 6: Median and 90th percentile interaction latencies as measured in Crux versus a single, global **memcached** instance.

mediately after its corresponding *Put*, we take the sum of these two latencies as “interaction latency” in the sense that this is the shortest possible delay with which a value could *in principle* propagate from u to v via this *Put/Get* pair, supposing v launched its *Get* at “exactly the right moment” to obtain the value of u ’s *Put*, but no earlier.

The collection of PlanetLab servers for these experiments have pairwise latencies ranging from 66 microseconds to 449 milliseconds. For simplicity, the servers in the experiment also act as the clients, initiating all *Put* and *Get* operations. In the presence of an occasional node failure, we discard the interaction from the results.

Figure 6 shows the interaction latency as defined above when using Crux versus a single **memcached** instance. For visual clarity, we group all interactions into one of thirteen “buckets” and then plot the median and 90th percentile latency for each bucket. Additionally, the graph plots the $y = x$ line as a reference for the expected lower bound on interaction latency.

The graph shows that *Put/Get* interaction latency in Crux closely approximates real internode network latency. For extremely close nodes (e.g., real internode latency $\ll 1$ ms), Crux appears to hit a bottleneck beyond the raw network latency, perhaps in **memcached** itself. However, the median Crux latency of roughly 1ms for these nearby nodes represents nearly a 3 orders of magnitude improvement on the median **memcached** latency. In a comparative baseline configuration with a single **memcached** instance distributed across the whole network (blue lines in Figure 6), the observed interaction latency unsurprisingly remains roughly constant, dominated by the graph’s global delay diameter and largely independent of distance between interacting nodes.

The CDF in Figure 7 shows the total number of *Put* and *Get* operations serviced by all nodes, in a single global **memcached** instance versus a Crux deployment using either inclusive or exclusive rings. The near-uniform load on all nodes in the baseline **memcached** in-




Figure 7: CDF of the total **memcached** operations serviced by a node participating in a single global instance versus using Crux with inclusive and exclusive rings.

stance indicates the consistent hashing scheme successfully distributes work. Both Crux variants incur significantly more overhead than the baseline case of a single **memcached** instance, and Crux’s non-uniform level assignment causes a somewhat wider variance in total operations each node services, though load remains generally predictable with no severe outliers. The order of magnitude increase in total operations is expected given the additional \mathcal{A} instance memberships for each node, and represents the main cost of preserving locality with Crux.

We make no claim that Crux is suitable in every distributed service or application scenario. Crux’s overhead may or may not be acceptable depending on the application, deployment requirements, and hardware, and especially the relative importance of locality preservation versus the costs of handling the additional load. Crux is likely to be more suitable for applications in which it is important to localize latency but request load is relatively low-volume and not a primary cost factor, such as in control-plane or event notification services for example. Section 3.4 briefly describes some potential optimizations for reducing these costs, but we leave these goals to future work.

5.3 Locality-Preserving DHT

We now evaluate Crux applied to the **bamboo** [4] open-source distributed hash table (DHT), comparing a Crux-transformed deployment against the baseline of a single network-wide **bamboo** instance. As in **memcached**, our goal is to exploit locality in DHT lookups by maintaining multiple DHT instances so that a node v ’s *Get* operation succeeds within the closest DHT instance that it shares with node u , the initiator of the *Put*. This experiment uses the same input graph and random key selection workload as in the previous experiment. Additionally, these experiments record Crux’s interaction latency in the same way: we observe all *Put* latencies for a given




Figure 8: Median and 90th percentile interaction latencies as measured in Crux versus a single, global DHT instance. The red line shows the lower bound on internode latency for the underlying DHT implementation.

key and then add the latency of the quickest *Get* lookup to the *Put* latency for the particular shared DHT instance.

Figure 8 shows the median and 90th percentile latencies for interaction via DHT *Put* and *Get* operations, using Crux versus a single global **bamboo** instance. The red line in the graph shows the minimum interaction latency we experimentally found to be achievable across any pair of nodes; this appears to be the point at which delays other than network latencies begin to dominate **bamboo**’s performance. To identify this lower bound, we created a simple 2-node DHT instance using nodes connected to the same switch with a network propagation delay of just 66 microseconds. We then measured the latency of a *Put* and *Get* interaction between these nodes, which consistently took roughly 10ms.

The graph shows promising gains at the median latency for nearby nodes in Crux when compared to a single DHT instance. As expected, the variance in the median and 90th percentile latencies for the single global DHT instance is quite small, once again showing that the expected delay for a non-locality-preserved algorithm is a function of the graph’s delay diameter. Crux shows a definite trend towards the $y = x$ line for real internode distances above approximately 10ms, but in general, the multi-hop structure of the DHT dampens the gains with Crux. While we expect this, the tightness with which Crux tracks the implementation-specific lower bound on internode latency suggests that a more efficient DHT implementation might increase Crux’s benefits further.

5.4 Locality-Preserving Pub/Sub

We finally evaluate Crux’s ability to reduce interaction latency in a publish/subscribe (pub/sub) distributed system based on **redis** [25]. Unlike the **memcached** and **bamboo** examples, where we defined interaction latency as the sum of distinct *Put* and *Get* operation latencies, a pub/sub interaction is a single contiguous process: an




Figure 9: Median and 90th percentile interaction latencies as measured in Crux versus a single, global pub/sub instance.

initiating node publishes a message to a named channel, at which point **redis** proactively forwards the message to any nodes subscribed to that channel.

This experiment uses the same input graph and random key selection workload as in the previous experiments, but changes the latency measurement method to account for contiguous pub/sub interactions. Crux deploys a set of independent **redis** servers for each \mathcal{A} instance, and clients use a consistent hash on the channel name to choose the **redis** server responsible for that channel. The experiment effectively implements a simple “ping” via pub/sub: two clients u and v each subscribe to a common channel, then u publishes a message to that channel, which v sees and responds to via another message on the same channel. Since this process effectively yields two pub/sub interactions, we divide the total time measured by u in half to yield interaction latency.

Figure 9 shows the median and 90th percentile measured interaction latencies for pub/sub interactions using Crux, again versus a single global pub/sub instance. As in **memcached**, the pub/sub experiment shows the greatest gains for nearby nodes, recording median latencies close to 1ms compared to **redis**’ normal median latency of over 500ms, orders of magnitude improvement over the baseline. Furthermore, pub/sub does not exhibit the bottleneck lower bound on interaction latency we observed with **bamboo**. As in all of the experiments, the median and 90th percentile latencies for the single global instance of **redis** hold steady around 500ms, remaining a function of the graph’s delay diameter and independent of internode network latency.

6 Related Work

The construction of multiple \mathcal{A} instances in Crux builds directly on landmark [30] and compact routing schemes [28, 29], which aim to store less routing information in exchange for longer routes. Techniques such as IP Anycast [2] and geolocation [13] or GeoDNS [11]

attempt to increase locality by connecting a client to the physically closest server, but these schemes work at the network level without any understanding or knowledge about the application running atop it.

Distributed storage systems [6, 27] and CDNs [1, 10] often employ data replication [12] and migration [5, 34] to reduce client interaction latency, with varying consistency models [16]. Some distributed hash table (DHT) schemes, such as Kademlia [20], include provisions attempting to improve locality. Many systems are optimized for certain workloads [6] or specific algorithms, such as Coral’s use of hierarchical routing in DHTs [10].

Like Crux, Piccolo [23] exploits locality during distributed computation, but it represents a new programming model, whereas Crux offers a “black box” transformation applicable to existing software. Canal [31] recently used landmark techniques for the different purpose of increasing the scalability of Sybil attack resistant reputation systems. Brocade [33] uses the concept of “supernodes” to introduce hierarchy and improve routing in DHTs, but selects supernodes via heuristics and makes no quantifiable locality preservation guarantees.

7 Conclusion

Crux introduces a general framework to create locality-preserving deployments of existing distributed systems. Building on ideas embodied in compact graph processing schemes, Crux bounds the latency of nodes interacting via a distributed service to be proportional to the network delay between those nodes, by deploying multiple instances of the underlying distributed algorithm. Experiments with an unoptimized prototype indicate that Crux can achieve orders of magnitude better latency when nearby nodes interact via the system. We anticipate there are many ways to develop and optimize Crux further, in both generic and algorithm-specific ways.

References

- [1] Akamai technologies.
<http://www.akamai.com>.
- [2] J. Abley and K. Lindqvist. Operation of anycast services, December 2006. RFC 4786.
- [3] Amazon Web Services. Amazon elastic compute cloud (Amazon EC2).
<http://aws.amazon.com/ec2/>.
- [4] The bamboo distributed hash table.
<http://bamboo-dht.org/>.
- [5] Brad Calder, Ju Wang, Aaron Ogus, Niranjani Nilakantan, Arild Skjolsvold, Sam McKelvie, Yikang Xu, Shashwat Srivastav, Jiesheng Wu, Huseyin Simitci, Jaidev Haridas, Chakravarthy Uddaraju, Hemal Khatri, Andrew Edwards, Vaman Bedekar, Shane Mainali, Rafay Abbasi, Arpit Agarwal, Mian Fahim ul Haq, Muhammad Ikram ul Haq, Deepali Bhardwaj, Sowmya Dayanand, Anitha Adusumilli, Marvin McNett, Sriram Sankaran, Kavitha Manivannan, and Leonidas Rigas. Windows azure storage: A highly available cloud storage service with strong consistency. In *23rd ACM Symposium on Operating Systems Principles (SOSP)*, SOSP ’11, 2011.
- [6] Fay Chang et al. Bigtable: A distributed storage system for structured data. In *7th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, November 2006.
- [7] Brent Chun et al. PlanetLab: An overlay testbed for broad-coverage services. In *ACM Computer Communications Review*, July 2003.
- [8] Al Danial. Counting Lines of Code, ver. 1.53.
<http://cloc.sourceforge.net/>.
- [9] Brad Fitzpatrick. Distributed caching with memcached. *Linux Journal*, August 2004.
<http://www.linuxjournal.com/article/7451>.
- [10] Michael J. Freedman, Eric Freudenthal, and David Mazières. Democratizing content publication with Coral. In *1st USENIX Symposium on Networked System Design and Implementation (NSDI)*, March 2004.
- [11] GeoDNS.
<http://www.caraytech.com/geodns/>.
- [12] Andreas Haeberlen, Alan Mislove, and Peter Druschel. Glacier: Highly durable, decentralized storage despite massive correlated failures. In *2nd Symposium on Networked Systems Design and Implementation (NSDI)*, May 2005.
- [13] Ethan Katz-Bassett, John P. John, Arvind Krishnamurthy, David Wetherall, Thomas Anderson, and Yatin Chawathe. Towards IP geolocation using delay and topology measurements. In *Internet Measurement Conference (IMC)*, October 2006.
- [14] Dmitri Krioukov, Kevin Fall, and Xiaowei Yang. Compact routing on Internet-like graphs. In *IEEE INFOCOM (INFOCOM)*, 2004.
- [15] Dmitri Krioukov, kc claffy, Kevin Fall, and Arthur Brady. On compact routing for the Internet. *Computer Communications Review*, 37(3):41–52, 2007.
- [16] Cheng Li, Daniel Porto, Allen Clement, Johannes Gehrke, Nuno Preguica, and Rodrigo Rodrigues. Making geo-replicated systems fast as possible, consistent when necessary. In *10th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, October 2012.
- [17] Cristian Lumezanu, Randy Baden, Neil Spring, and Bobby Bhattacharjee. Triangle inequality and routing policy violations in the Internet. In *Passive and Active Measurement Conference (PAM)*, April 2009.
- [18] Qin Lv, Pei Cao, Edith Cohen, Kai Li, and Scott Shenker. In *16th International Conference on Supercomputing (ICS)*, June 2002.

- [19] Harsha V. Madhyastha, Tomas Isdal, Michael Poatek, Colin Dixon, Tomas E. Anderson, Vrvind Krishnamurthy, and Arun Venkataramani. iPlane: an information plane for distributed services. In *OSDI*, 2006.
- [20] Petar Maymounkov and David Mazières. Kademlia: A peer-to-peer information system based on the XOR metric. In *1st International Workshop on Peer-to-Peer Systems (IPTPS)*, March 2002.
- [21] Nick McKeown, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Scott Shenker, and Jonathan Turner. OpenFlow: Enabling innovation in campus networks. *ACM SIGCOMM Computer Communication Review*, 38(2):69–74, 2008.
- [22] Charles E. Perkins, Elizabeth M. Belding-Royer, and Samir Das. Ad hoc on-demand distance vector (AODV) routing, January 2002. Internet-Draft (Work in Progress).
- [23] Russell Power and Jinyang Li. Piccolo: Building fast, distributed programs with partitioned tables. In *9th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, October 2010.
- [24] William Pugh. Skip lists: a probabilistic alternative to balanced trees. *Communications of the ACM*, 33(6):668–676, 1990.
- [25] The redis pub sub. <http://redis.io/topics/pubsub>.
- [26] A. Rowstron and P. Druschel. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In *International Conference on Distributed Systems Platforms (Middleware)*, 2001.
- [27] Ion Stoica et al. Chord: A scalable peer-to-peer lookup service for Internet applications. In *ACM SIGCOMM*, August 2001.
- [28] Mikkel Thorup and Uri Zwick. Approximate distance oracles. In *ACM Symposium on Theory of Computing*, pages 183–192, 2001.
- [29] Mikkel Thorup and Uri Zwick. Compact routing schemes. In *ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 1–10, New York, NY, USA, 2001. ACM Press.
- [30] Paul Francis Tsuchiya. The Landmark hierarchy: A new hierarchy for routing in very large networks. In *ACM SIGCOMM*, pages 35–42, August 1988.
- [31] Bimal Viswanath, Mainack Mondal, Krishna P. Gummadi, Alan Mislove, and Ansley Post. Canal: Scaling social network-based sybil tolerance schemes. In *EuroSys Workshop on Social Network Systems (SNS)*, April 2012.
- [32] Steve Yegge. Stevey’s google platforms rant, October 2011. <https://plus.google.com/+RipRowan/posts/eVeouesvaVX>.
- [33] Ben Y. Zhao, Yitao Duan, Ling Huang, Anthony D. Joseph, and John Kubiawicz. Brocade: Landmark routing on overlay networks. In *1st International Workshop on Peer-to-Peer Systems*, March 2002.
- [34] Ben Y. Zhao, Ling Huang, Jeremy Stribling, Sean C. Rhea, Anthony D. Joseph, and John D. Kubiawicz. Tapestry: A resilient global-scale overlay for service deployment. *IEEE Journal on Selected Areas in Communications*, 22(1):41–53, January 2004.